Performance Studies for FDDI Connectivity

Michelle R. Koblas
Jonathan Hahn

Numerical Aerodynamic Simulation Systems Division
NASA Ames Research Center, M/S 258-6
Moffett Field, CA 94035-1000

September 22, 1992

Abstract

This paper analyzes the possibility of replacing Network Systems Corporation (NSC) N130 adapters providing FDDI (Fiber Distributed Data Interface) ring connectivity for Crays and an Amdahl 5880 with a 1000 series Hub from Ultra Networking Technologies, Inc. in conjunction with a Cisco Systems router. We look at the experimental results from two testbed setups with these network architectures and examine the throughput achieved in terms of transfer rates and the problems encountered. The results show that the Ultra and Cisco configuration does not provide any performance enhancement over the existing NSC equipment and, as such, it is not recommended to replace the existing equipment.

1.0 Introduction

The Numerical Aerodynamic Simulation (NAS) facility currently uses equipment from Network Systems Corporation (NSC) to connect its Cray YMP, Cray 2 and Amdahl to an FDDI (Fiber Distributed Data Interface) ring. It was speculated that utilizing equipment from Ultra Networking Technologies, Inc. in conjunction with a Cisco Systems, Inc. router to connect these machines to the FDDI ring might provide faster throughput. This paper presents the results of testing the performance of the two network architectures for connecting the Crays and the Amdahl to an FDDI ring.

2.0 Experimental Setup

To test the hypothesis that the combination of Ultra and Cisco equipment would provide better performance than the current NSC equipment for connecting the Crays and Amdahl to the FDDI ring, two testbeds were set up. The first testbed utilized NSC equipment and the second used the Ultra/Cisco architecture. This section explains the hardware and software used during testing.

2.1 Equipment

The following equipment was used for testing:

reynolds	Cray YMP running Unicos 6.1.6
sandbox	Amdahl 5880 running UTS 2.1
ogre	Silicon Graphics, Inc. (SGI) workstation (4d320VGX) running IRIX 4.0.3 FDDIXPress controller version 1.0 FDDIXPress system software version 2.0 Ultra VME adapter, 250Mbits/sec, with software version 4.0.16
igson	SGI workstation 4d440VGX running IRIX 4.0.1 FDDIXPress controller version 1.0 FDDIXPress system software version 2.0
cham100	Tekelec ChameLAN 100s FDDI Analyzer, running monitoring software version 3.0
ron & *judith*	Network Systems Corporation N130 Adapters running CDA version 1.0
cisco	Cisco AGS+ router running version 9.0 FDDI controller type 4.0, microcode version 1.0 C-Bus controller type 3.0, microcode version 2.0 Ultra Cisco adapter
Hub 10	Ultra hub 1000 Network Manager software release 2.50.27 Adapter code version 3.00.24
wellfleet	Wellfleet LN Model 2000 router running version 5.74

Although the Cray 2 was also available, only the Cray YMP (*reynolds*) was used for performance measuring on the assumption that the Crays would have similar performance comparisons across the two network architectures.

2.2 Hardware Configurations

The first testbed, shown in Figure 1, was designed to utilize existing NSC equipment. For this test, two NSC N130 adapters were removed from the production FDDI ring and placed in a test ring along with a Wellfleet router and two SGI workstations (*ogre* and *igson*). *Reynolds* was connected to an NSC adapter (*judith*) via a Cray low speed channel (herein referred to as 'LSC') and *sandbox* connected to the other NSC adapter (*ron*) via a pair of striped Block Multiplexer Channels (BMC). Tests were run between *reynolds*, *sandbox* and *ogre*, as well as some metering between the workstations on the FDDI ring.

FIGURE 1. NSC Configuration

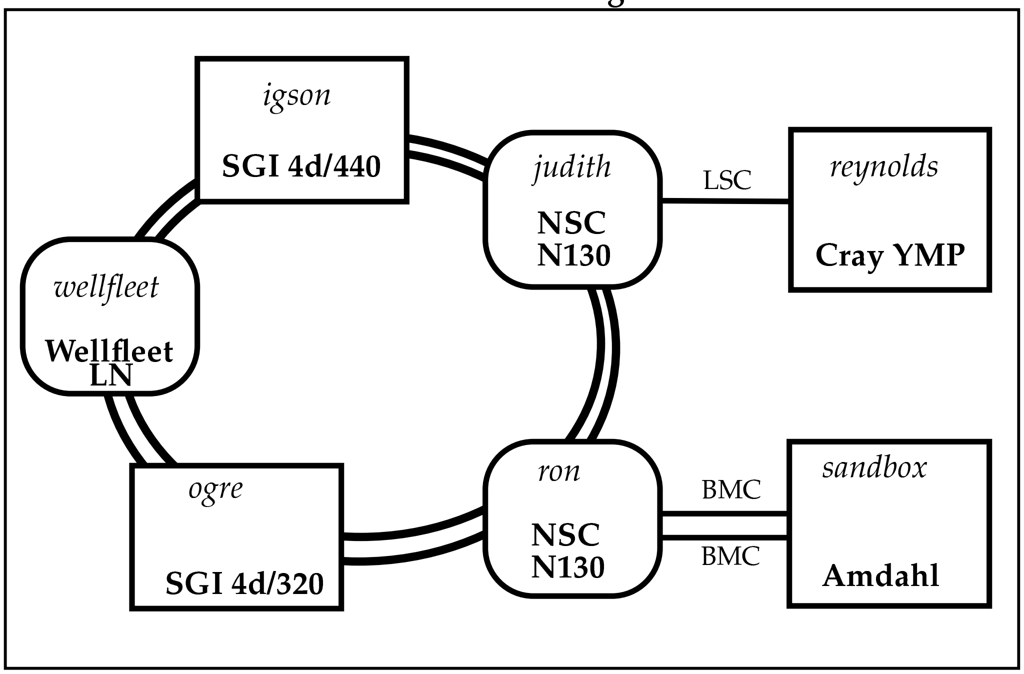

In the second setup, shown in Figure 2, an Ultra Hub 1000 was used in conjunction with a Cisco router to provide FDDI connectivity for *reynolds* and *sandbox*. S*andbox* was connected to the hub via a pair of striped BMC channels. The connection for *reynolds* was tested using both a Cray LSC as well as a Cray high speed channel (HSX) to connect to the hub. The Ultra adapters for both the Cray and Amdahl were run in "hybrid"[1] mode, while both the Cisco router and *ogre* were connected to the Ultra hub with their software running in "new" mode[2]. The physical connections for the Cisco router and *ogre* were paired coaxial cables from a multiplexer board in the

1. The Ultra Networking Technologies, Inc. equipment is in a period of transition. Originally, the network ran internally using an OSI protocol, but this is being changed to utilize TCP/IP. This has resulted in three modes of operation which will be referred to as "old", "new", and "hybrid" modes. "Old" mode uses OSI, "new" mode uses TCP/IP, and "hybrid" mode can communicate with either of the other modes and is in use only on channel-based machines (Cray and Amdahl).

2. The Cisco router had to be in "new" mode to connect to the ultra hub.

hub to an Ultra VME bus adapter for the SGI workstation and a HSSI (high speed serial interface) connection on the Cisco router.

FIGURE 2. Ultra and Cisco Configuration

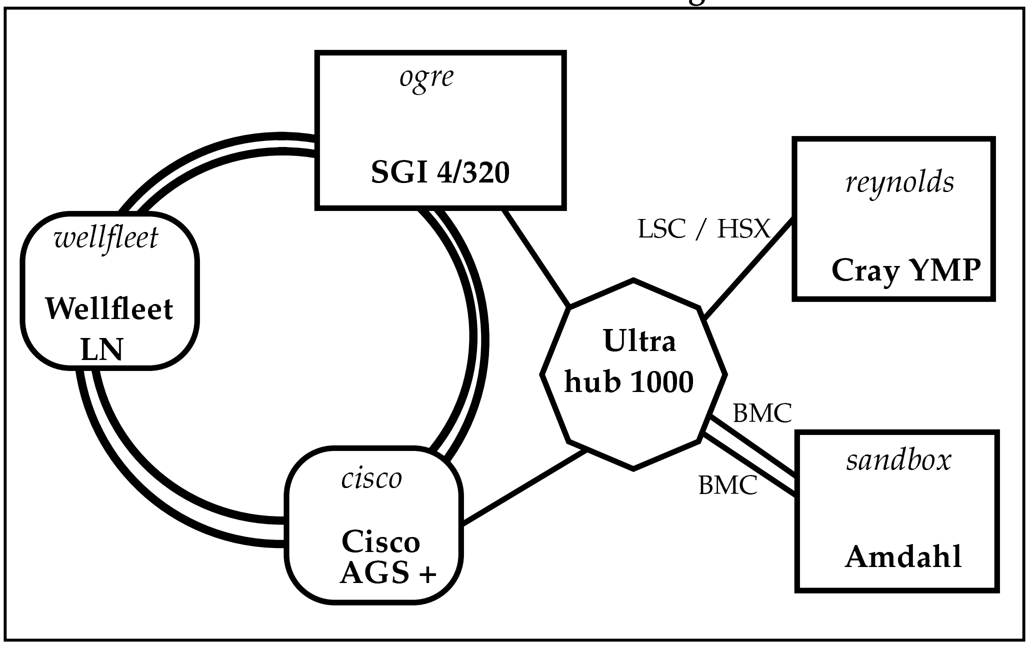

The Wellfleet router shown in both configurations was not used for any throughput tests, but resided on the experimental FDDI rings in order to make the testbeds more closely resemble the configuration in which the equipment would run in the NAS network. Additionally, a Tekelec ChameLAN 100s FDDI analyzer, *cham100*, is also on the testbed FDDI rings, but not shown in the figures. It is a passive FDDI monitor and does not affect traffic on the ring.

2.3 Interface Speeds

Since the significant portion of this experiment relies on the transfer rates between hosts for making conclusions, it is important to know the speeds at which each of the interfaces or network types are theoretically capable of operating. Table 1 contains this information.

TABLE 1.

Interface/Network	Theoretical Bandwidth
FDDI	100 Mbits/second
Cray LSC	100 Mbits/second
Cray HSX	800 Mbits/second
BMC	36 Mbits/second
Cisco Ultra interface	125 Mbits/second
SGI Ultra VME	250 Mbits/second
Ultra Hub backplane	1 Gigabit/second

The number shown for the BMC interface is for a single channel and two striped BMCs are used in the testbeds, which makes the theoretical speed for the pair 72 Mbits/second. The Ultra hub backplane speed is the rate which data can be passed between adapters in the Ultra hub, so it can be considered the speed of the UltraNet[3].

2.4 Software Tools

All test results reported in this paper were attained using the *tsock* program in source/sink mode sending TCP packets, with user buffers ranging from 64 bytes to 64 Kilobytes. The tests were performed from memory to memory in a series of three transfers at each buffer size. The buffers for the first testbed only range up to 8 Kilobytes, as it was not clear at that time that larger buffers would be interesting. Some of the resulting graphs do not level out at the highest buffer size, so it is possible that larger buffers might produce higher transfer speeds than reported here through the NSC equipment.

2.5 Test Procedure

Both testbeds (NSC and Ultra/Cisco) were used to test throughput measurements from *reynolds* and *sandbox* to *ogre* and vice versa. Additionally, measurements were taken between machines over only the UltraNet in the second testbed and between two SGIs (*ogre* and *igson*) on the FDDI ring in the first testbed to be used as baseline figures for these networks.

3.0 Results

3.1 Graph Format

The results presented below are in the form of graphs, with the lines labelled by source and destination hosts, media used and, where necessary, interface in the following format:

‘<source>_<destination>.<media>[<interface>][.<mode>]’

<Media> refers to these network paths:

C = Cisco/Ultra network path

N = NSC equipment only

U = UltraNet only

F = FDDI ring only

3. “UltraNet” is used to refer to the path between hosts which only traverses the proprietary network of Ultra Networking Technologies, Inc. equipment.

<Interface> is only valid on the Cray YMP and is either “L”, for the Low Speed Channel, or “H” to represent using the High Speed channel.

<Mode> can be either “old”, “new”, or “hybrid”. This notation is only used for the graph of *reynolds’* loopback testing and refers to the mode of operation of *reynolds’* Ultra adapter.

3.2 NSC equipment testbed

The testbed of NSC equipment, pictured in Figure 1 above, was used for testing between *reynolds* and *ogre*, as well as *sandbox* and *ogre*. For each test in this suite, transfers using from 64 byte to 8196 byte user buffers were repeated three times and the graph shown in Figure 3 reports the best rates at each buffer size for the various pairs of hosts.

FIGURE 3. NSC Transfer Rates

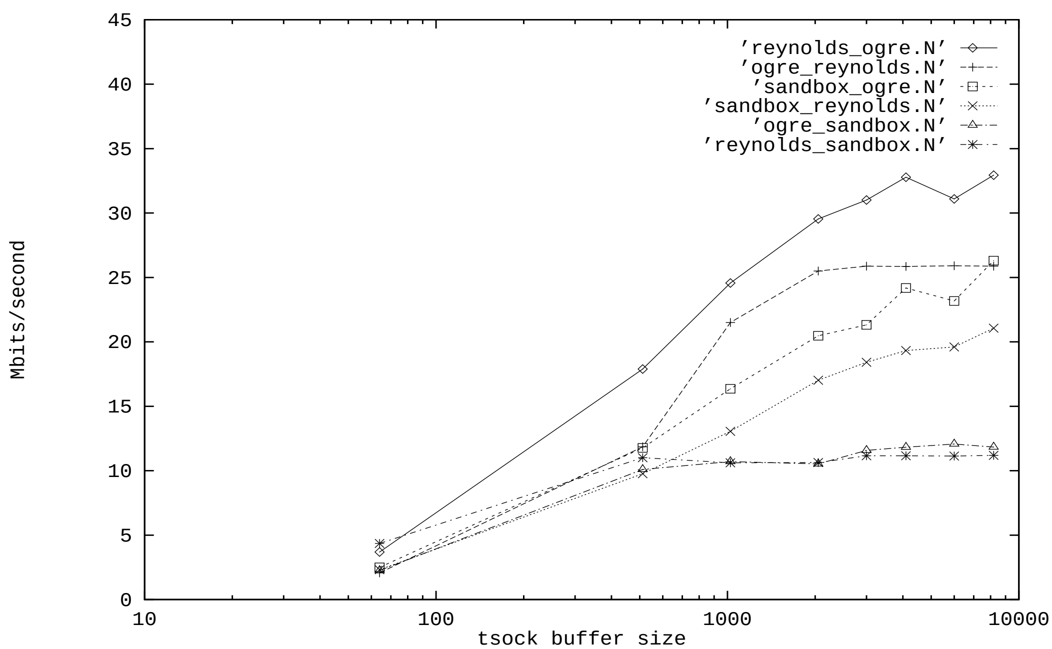

As can be seen from the graph, transfers to and from *reynolds* through the NSC equipment produced the highest transfer rates, reaching a peak of 32.9 Mbits/second when using 8196 byte user buffers. This was the expected result, considering the Cray YMP was the fastest host used. It is interesting to note, however, that although *sandbox* produced the second fastest rates sending data, the receiving speed seems to reach a threshold at approximately 12 Mbits/second. The highest transfer rates seen during this test into the Amdahl were 11.8 Mbits/second, at 4096 byte buffers. Later tests, using the Ultra/Cisco equipment, produced a similar plateau, suggesting something wrong with the receive portion of this interface.

The NSC testbed was also used to generate transfer rates between the two SGI workstations on the FDDI ring, in order to determine what speeds these machines (and the FDDI ring in general) were capable of reaching. The results of this test showed that between the workstations (*igson* and *ogre*), transfer rates of up to 40.0 Mbits/second were attainable with ogre transmitting 8196 byte user buffers. Thus, slower transmission rates between *ogre* and hosts behind the NSC adapters can be attributed to either the receiving host or the NSC equipment.

3.3 Ultra and Cisco equipment testbed

There were many problems using the Ultra/Cisco equipment testbed (shown in Figure 2 on page 4), mostly related to the Ultra hardware and software. Testing was begun between hosts using only the UltraNet (i.e., without involving the Cisco router or FDDI ring). This set of tests was designed to determine the rates of transfers between hosts on the UltraNet to compare with rates through the Cisco router, in order to determine if the router created any type of bottleneck. All tests in this testbed were run between hosts using the Ultra "host stack"[4].

At first, the Cray LSC was used to connect *reynolds* to the UltraNet, since this is the same type of interface which the Cray uses to connect to the NSC adapter. However, there was a problem establishing a connection to the Cray from *ogre* when attempting to get transfer rates over only the UltraNet. This problem manifested itself as an extremely slow connection to the Cray interface which degraded as more packets were sent over it.Occasionally no connection could be established at all. This was attributed to a problem in the SGI workstation's Ultra VME adapter microcode running in "new" mode when attempting to communicate with channel-based adapters. New microcode was installed for the VME adapter which was believed to fix the problem. Unfortunately, similar symptoms appeared again when eventually the Cray HSX interface was used (in an attempt to get better numbers) and tests in that direction (*ogre* to *reynolds*) which traverse only the UltraNet are not included in the results.

Throughput tests between hosts using only the UltraNet produced numbers up to 11.8 Mbits/second from *reynolds* to *ogre*, through the HSX adapter and up to 5.7 Mbits/second from *sandbox* to *ogre* with both the HSX and BMC adapter in "hybrid" mode. These tests were expected to produce much better throughput numbers than were received, due to a previous report [3], however, upon further investigation, it appears that there is a significant difference in speeds between Ultra adapters running in "hybrid" versus "old" modes. To verify this, a quick series of loopback

4. The UltraNet consists of two logical networks operating on the same physical layer. The two methods of accessing this network are called "host stack", for transfers using the host's kernel TCP/IP stack, and "native" for transfers which bypass TCP/IP in the kernel. See [2] for a more detailed description.

tests from *reynolds* to itself was tried with the Cray HSX adapter in place. The graph in Figure 4 includes the results of these tests. One possible explanation of this difference is the theory that the adapter must do more work converting between the various modes of operation than if it was strictly transferring data. It is possible that transfer rates will once again go up when the hosts are capable of communicating with the adapters running in “new” mode. Although this test was not attempted with the LSC adapter or the *sandbox* connection, it is felt that these results are important in understanding the outcome of this experiment.

FIGURE 4. Cray HSX Loopback in “Hybrid” vs. “Old” Mode

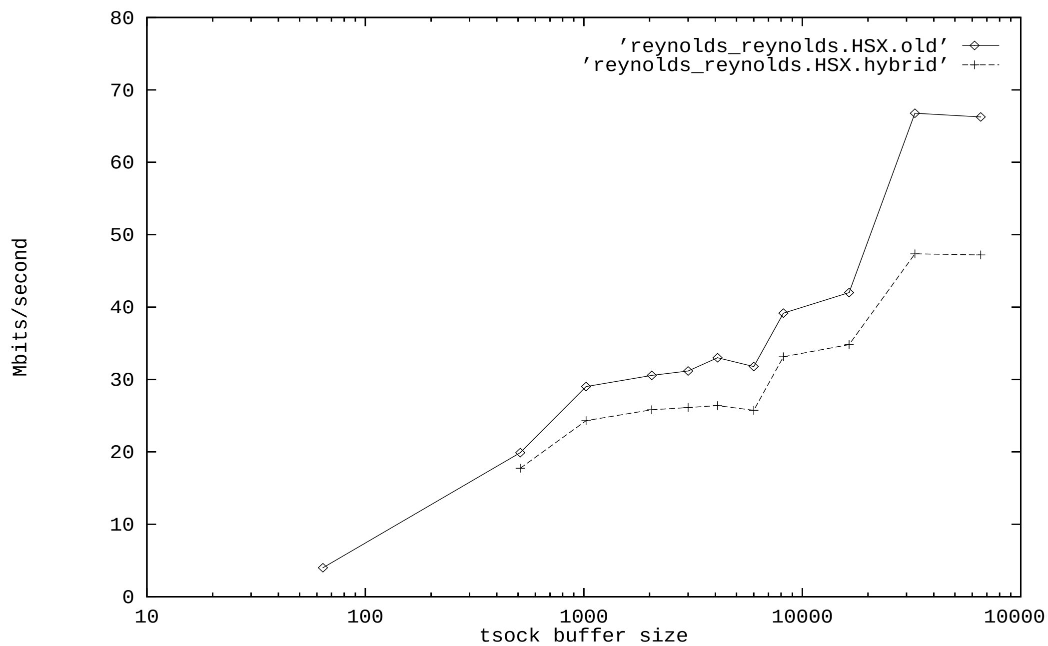

An additional problem with this testbed was poor transfer rates through the Cisco router due to misconfigured parameters on the Cisco router. The router was configured such that it was not using cBus caching and had too large an MTU (Maximum Transmission Unit) size on the router’s UltraNet interface. The graph in Figure 5 shows the comparison between transfers over the UltraNet versus transfers which also go through the Cisco router between *ogre* and *reynolds*, using the HSX interface.

FIGURE 5. Ultra vs. Ultra/Cisco Transfers Through the Cray HSX

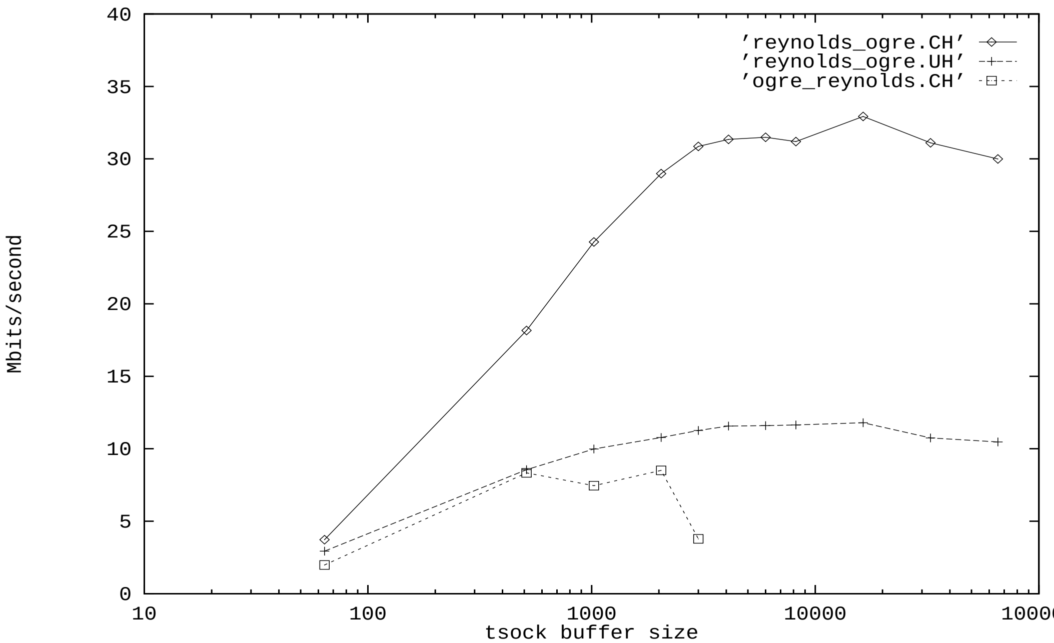

As can be seen from the above graph, the transfer rates received are somewhat confusing. One would expect that additional equipment between the two hosts would produce slower speeds, however, the addition of the Cisco router and FDDI ring between *reynolds* and *ogre* appears to produce better results! This phenomenon is not limited to the *reynolds* to *ogre* transfers. As can be seen on a later graph, *sandbox* to *ogre* transfers are also faster through the Cisco router and FDDI than through only the UltraNet. This suggests some problem in the Ultra hardware/software which requires investigation. More specifically, it appears that there is a major problem in the SGI's UltraNet connection, either in the hardware or software when traversing the "host stack" path. The "native" UltraNet path of this interface was beyond the scope of this experiment, although it could shed much light on whether it is a problem in the hardware or software.

Also, notice the 'ogre_reynolds.CH' graph line. This transfer eventually began to produce such poor transfer rates that it became impossible to continue testing above 4 Kilobyte sized user buffers. Unfortunately, testbed time was not sufficient to warrant an in-depth investigation of this problem.

4.0 Conclusion

Although the Ultra and Cisco equipment was expected to produce much higher transfer rates than the current NSC equipment at NAS, the graphs in Figure 6 and Figure 7 show this to not be the case. Transfers between hosts using the NSC equipment in the first testbed produced significantly better rates than the corresponding transfers through Ultra/Cisco equipment in the second testbed.

FIGURE 6. Transfer Rates between *Reynolds* and *Ogre*

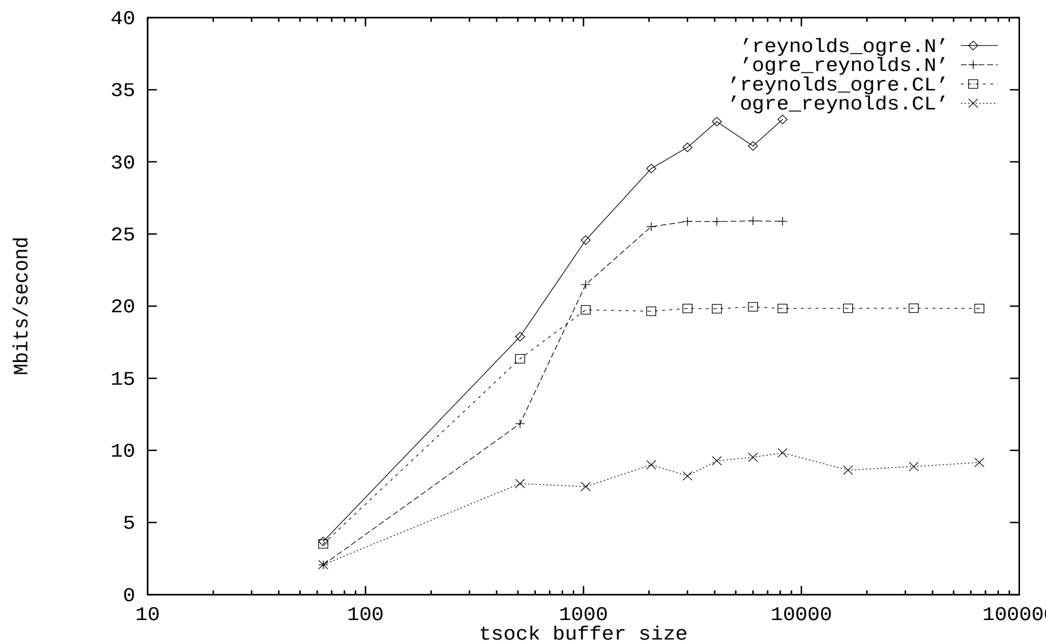

FIGURE 7. Transfer Rates between *Sandbox* and *Ogre*

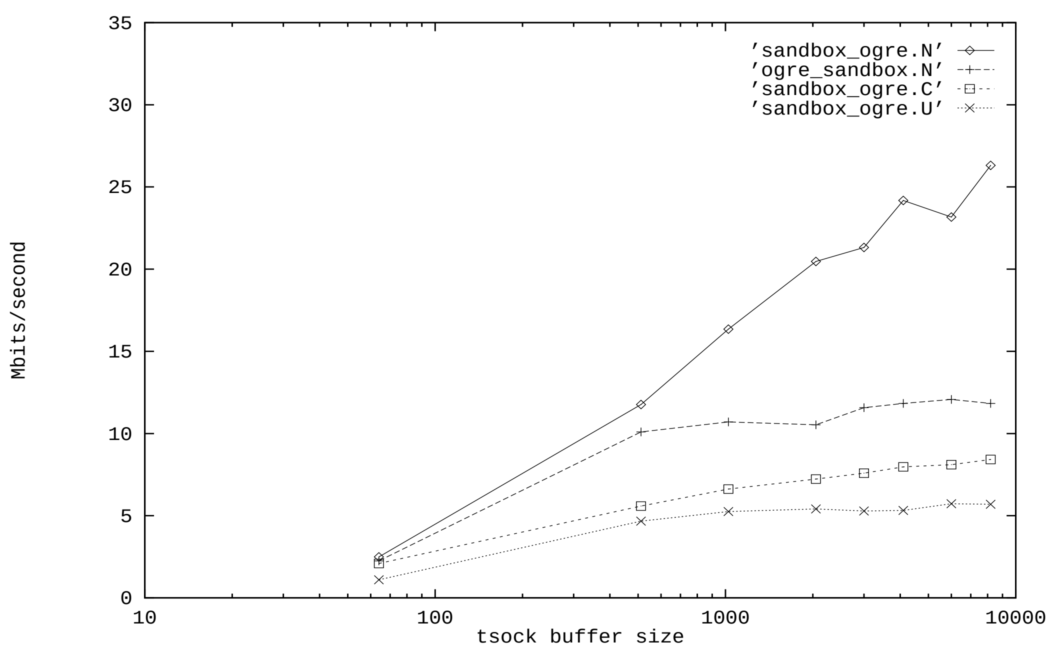

Some of the difference from the original expectation of the Ultra equipment speed could be attributed to the slower rates seen over the UltraNet with the adapters operating in "hybrid" mode versus the rates in "old" mode, however, nothing can be proven conclusively. It is obvious that there are some issues with the Ultra hardware/software which need to be investigated further. It would be an interesting addition to this experiment to arrange a similar Ultra/Cisco testbed when the channel-based hosts (Amdahl and Cray, in this case) and their adapters are capable of running totally in Ultra "new" mode; however it is unclear when this software will be available.

Thus, at this time, based on the comparison of transfer rates exhibited between hosts through the current NSC equipment and through the Ultra hub and Cisco router, it is not recommended that Ultra and Cisco equipment be purchased to replace the existing NSC adapters providing Cray and Amdahl connectivity to the FDDI ring.

5.0 References

[1] *Network Addressing Manual*, Ultra Network Technologies, Inc., Part number 06-0020-001, Revision B.

[2] *Ultra Product Overview,* Ultra Network Technologies, Inc., Part number 06-0017-001, Revision A.

[3] *Very High Speed Network Prototype Development*, *Task 2.1: Measurement of Effective Transfer Rates,* Marke Clinger, October 25, 1989, Ultra Network Technologies, Inc.